

An **Appendix** including replacement drawing figures is attached following **page 13** of this paper.

Amendments to the Specification:

*Please **replace paragraph [0007]** with the following amended paragraph:*

[0007] Disclosed is a method and system of replicating data using a recovery data change log. According **to** one embodiment of the present invention, a copy of a data change log is maintained at a data recovery node, the data change log being associated with a primary data volume of a primary node, a failure of the primary data volume is detected, and a secondary data volume of a secondary node is updated using the copy of the data change log in response to the detection.

*Please **replace paragraph [0027]** with the following amended paragraph:*

[0027] During asynchronous replication, an application write operation completes as soon as the write or “update” is logged at the primary data change log 118. Consequently, transmission to and storage at secondary node(s) is concurrent with **continues continued** application execution at the primary node. Asynchronous replication typically reduces application response time and allows applications to continue execution during momentary network outages or overloads.

*Please **replace paragraph [0031]** with the following amended paragraph:*

[0031] In the illustrated embodiment of Fig. 1B, secondary data volume 120 is updated by copying one or more entries from recovery data change log 124 to secondary data change log 122 and updating secondary data volume ~~129~~ **120** using secondary data change log 122. In one embodiment of the present invention, recovery data change log 124 includes a plurality of entries, each of which is associated with a requested write operation. The status of each write operation is determined such that, according to one embodiment, only those entries of recovery data change log 124 corresponding to

incomplete write operations (e.g., those for which a final replication acknowledgment has not been received at the primary node) are copied to secondary data change log 122.

*Please **replace** paragraph [0032] with the following amended paragraph:*

[0032] Once the appropriate entries of recovery data change log 124 have been identified, the identified entry(ies) are copied from the recovery data change log 124 to the secondary data change log 122 using volume replicator component(s) of volume manager(s) and/or volume replicator(s) 114c and 114b, respectively as represented by line 132. In alternative embodiments of the present invention however, any of a number of other data transport techniques or means are used to perform the described copying. In the illustrated embodiment, the described copying/replication is performed using recovery module ~~124~~ 126. In alternative embodiments of the present invention, recovery module ~~124~~ 126 may be implemented as data processing system software, hardware, firmware, and/or a combination thereof.

*Please **replace** paragraph [0034] with the following amended paragraph:*

[0034] In the illustrated embodiment of Fig. 1C, secondary data volume 120 is updated by copying one or more blocks of data, each block of data including one or more entries (e.g., entries corresponding to incomplete write operations at primary node 102), from recovery data change log 124 to a staging data change log 134 over communications network 108 (e.g., a wide area network) as represented by line 136. The copied block(s) are applied in write-ordered sets to secondary data change volume 122 which is used to update secondary data volume ~~120~~ 120.

Please replace paragraph [0039] with the following amended paragraph:

[0039] Fig. 3 illustrates a process flow diagram of a replication and recovery process according to an embodiment of the present invention. In the illustrated process embodiment, a primary data volume of a primary node and a secondary data volume of a secondary node are initially synchronized (process block 302). Such ~~initially~~ initial synchronization may be implemented according to various embodiments of the present invention using data transfer from one node data processing system to another across a network, tape or other persistent backup and restore capabilities, or one or more snapshots or portions thereof. Following the initial synchronization of the described nodes and associated data volumes a determination is made whether or not a primary data volume failure has occurred (process block 304).